

PAPER

Quantum control landscape for a two-level system near the quantum speed limit

To cite this article: Martín Larocca *et al* 2018 *J. Phys. A: Math. Theor.* **51** 385305

View the [article online](#) for updates and enhancements.

Related content

- [Optimization search effort over open quantum system control landscapes](#)
Anand Oza, Alexander Pechen, Jason Dominy et al.
- [Enhancement of quantum speed limit time due to cooperative effects in multilevel systems](#)
P M Poggi, F C Lombardo and D A Wisniacki
- [Strong-field control landscapes of coherent electronic excitation](#)
Tim Bayer, Matthias Wollenhaupt and Thomas Baumert

Recent citations

- [Geometric quantum speed limits and short-time accessibility to unitary operations](#)
Pablo M. Poggi
- [Deep reinforcement learning for quantum gate control](#)
Zheng An and D. L. Zhou
- [Reinforcement learning for autonomous preparation of Floquet-engineered states: Inverting the quantum Kapitza oscillator](#)
Marin Bukov



IOP | ebooks™

Bringing you innovative digital publishing with leading voices to create your essential collection of books in STEM research.

Start exploring the collection - download the first chapter of every title for free.

Quantum control landscape for a two-level system near the quantum speed limit

Martín Larocca¹, Pablo M Poggi^{1,2}  and Diego A Wisniacki¹

¹ Departamento de Física ‘J. J. Giambiagi’ and IFIBA, FCEyN, Universidad de Buenos Aires, 1428 Buenos Aires, Argentina

² Center for Quantum Information and Control, University of New Mexico, Albuquerque, New Mexico 87131, United States of America

E-mail: larocca@df.uba.ar

Received 23 April 2018, revised 12 July 2018

Accepted for publication 27 July 2018

Published 21 August 2018



Abstract

The core problem in optimal control theory applied to quantum systems is to determine the temporal shape of an applied field in order to maximize the expected value of some physical observable. The complexity of this procedure is given by the structural and topological features of the quantum control landscape (QCL)—i.e. the functional which maps the control field into a given value of the observable. In this work, we analyze the rich structure of the QCL in the paradigmatic Landau–Zener two-level model, and focus in particular on characterizing the QCL when the total evolution time is severely constrained. By studying several features of the optimized solutions, such as their abundance, spatial distribution and fidelities, we are able to rationalize several geometrical and topological aspects of the QCL of this simple model and identify the effects produced by different types of constraint.

Keywords: quantum optimal control, control landscape, quantum speed limit

(Some figures may appear in colour only in the online journal)

1. Introduction

The development of new technologies based on quantum information processing is currently blossoming. Proposals for communication, computation and simulation protocols based on quantum mechanical effects [1–3] are nowadays being transformed into reality thanks to the extraordinary capabilities of such physical platforms as cold atoms in optical lattices, quantum dots and superconducting qubits [4–6]. To take full advantage of this, scientists rely on their growing ability to control physical systems in the quantum regime by using properly tailored external fields. In this context, optimization methods originally put forward in the late 1980s have proven to give robust control strategies [7, 8].

The typical problem to solve is to find the control field $\epsilon(t)$ which maximizes a certain objective functional $J[\epsilon]$, i.e. the probability of reaching a target state, for instance. Extensive application and study of these quantum optimal control techniques over the past few decades has given evidence of the benign features of what is commonly called the quantum control landscape (QCL)—that is, the functional dependence of the objective J on the field ϵ . In a seminal work [9], Rabitz *et al* showed that, under certain conditions, the QCL was devoid of sub-optimal local maxima (see also [10] for a prior related work). This result explained the extraordinary success of local optimization procedures. This remarkable result about the topology of QCLs has been intensively tested and studied over recent decades [11–13], and limitations are known to arise in various cases. For instance, local maxima or traps are expected to appear when the control problem has constraints [14–17]—due, for example, to the time-discretization of the fields or bandwidth and amplitude limitations imposed to them. Another interesting constraint is given by the evolution time of the system. Traps have been shown to exist in the vicinity of the minimal or quantum speed limit (QSL) time [18]³, and there have also been numerous reports of slowing down of optimization algorithms in that regime [19, 20]. However, systematic analysis on how exactly these constraints affect the control landscape have been limited as to now, and a joint assessment of multiple types of constraints is currently lacking in the literature.

In this work, we present a systematic analysis on the effects of controlling near the QSL (i.e. restricting the evolution time) on the structural and topological features of QCLs. In parallel, we also consider constraints given by the coarse-graining of temporal variable, which is inherent to any numerical implementation. This is, to the best of our knowledge, the first joint assessment of the interplay between these types of constraint on the features of the control landscape in a quantum system. Note that for a single control field, the optimization space has a dimension of N_{ts} , which is the number of time slots we use to discretize the temporal variable. Typically, N_{ts} may be of the order of 10^2 or 10^3 , and so it is not trivial to assess the global properties of $J[\epsilon]$, whose representation is given by a hypersurface in a $N_{ts} + 1$ dimensional space. We therefore propose a number of strategies to probe the QCL in order to obtain information about its features. By using random initial seeds, we explore a certain region of the parameter space, and using standard local optimization techniques, we arrive at optimized solutions. For these, we study (i) the distances between them, which allows us to probe the number and distribution of maxima in such region; (ii) their fidelities, which give us information about the emergence of traps due to the constraints imposed on the problem; and (iii) a structural parameter R defined in [21], which measures how straight the path is between the initial seed and the optimized solution. This parameter allows us to observe structural properties of the landscape.

We use as a testbed for our analysis a simple, yet paradigmatic model of a driven two-level quantum system which is described by the Landau–Zener Hamiltonian. This model has been widely applied in quantum physics, as it describes non-adiabatic transitions [22], Landau–Zener–Stuckelberg interferometry [23] and quantum phase transitions [24]. For this model, a related study explicitly showed that the control landscape is indeed devoid of traps for $N_{ts} \rightarrow \infty$ [15]. Here, we go beyond that result and characterize not only the topology of the landscape but also its geometrical structure as a function of both N_{ts} and the evolution time T . Although local maxima (traps) disappear in the limit of continuous field, global maxima are shown to exhibit an interesting two-clan structure in the vicinity of the quantum speed limit

³ Here we use ‘minimal control time’ and ‘QSL’ as synonyms for the shortest process duration with perfect control (fidelity equal to one). Other works consider QSL as a bound for the minimal evolution time between initial and final states.

(QSL). Moreover, the two families merge at the QSL, rendering only one (global) maximum for any landscape with $T < T_{\min}$. Regarding the geometry of the landscape, we analyze the straightness of the trajectories traversed by a pure gradient algorithm towards the maxima, by comparing the actual path length and the Euclidean distance between initial and optimized fields. These trajectories through control space are found to bend as T approaches T_{\min} , and a discontinuous jump is observed at $T = T_{\min}$ where every path from hundreds of random seeds reach the only global maximum in a perfect straight line. That is, the landscape is found to be trivial at the quantum speed limit.

In order to gain intuition about the different measures we propose, we perform a detailed study on the two-dimensional (2D) (i.e. $N_{\text{ts}} = 2$) control problem. In this case we can visualize the landscape as a surface in a three-dimensional space, and thus directly relate it to our numerical optimization results. Interestingly, we also find that many features already seen in this simple model carry over to higher-dimensional landscapes.

This paper is organized as follows. In section 2, we present the basics of optimal control theory and its application to the Landau–Zener two-level model. In section 3, we will present a ‘toy model’ in which the control field is discretized into just $N_{\text{ts}} = 2$ time steps. This will allow us to visualize the landscape directly, and thus will be helpful in designing strategies that allow us to probe its features in a more general setting. In section 4, we discuss such strategies and show results for general landscapes with $N_{\text{ts}} > 2$. Finally, in section 5, we present some concluding remarks.

2. Quantum optimal control and the Landau–Zener model

Consider the time evolution of an isolated driven quantum system described by the following Schrödinger equation

$$i \frac{d\hat{U}_t}{dt} = [\hat{H}_0 + \epsilon(t)\hat{H}_c] \hat{U}_t, \quad (1)$$

where \hat{U}_t is the unitary evolution operator of the system at time t , H_0 and H_c are the drift and control Hamiltonians respectively, and $\epsilon(t)$ symbolizes the control field. Note that we set $\hbar = 1$ from here on. Optimal control theory assesses the problem of deriving the shape of $\epsilon(t)$ that maximizes the value of a cost functional $J[\epsilon]$. For example, a typical goal of control tasks is to take a given initial state $|i\rangle$ to a desired target state $|f\rangle$ in a (fixed) time $t = T$ —that is, to obtain $U_T|i\rangle = |f\rangle$ (up to some global phase). Finding the fields that perform the desired task with the best possible accuracy is identical to locating the global maxima of the QCL $J[\epsilon]$, which, in this particular case, would simply take the form $J[\epsilon] = |\langle f|\hat{U}_T|i\rangle|^2$. We can look for such maxima by proposing an initial seed $\epsilon_{(0)}(t)$ for the field, and to update it iteratively by using information about the gradient of $J[\epsilon]$. This is the idea behind most of the optimal control methods which have been widely incorporated by quantum scientists in recent decades, such as Krotov [8, 25], GRAPE [26] and others. An algorithm of this type would generate a path through the landscape which connects $\epsilon_{(0)}$ to some optimal field $\epsilon_{(K)}(t)$, where K denotes the number of iterations. Note that while, ideally, we expect $J[\epsilon_{(K)}] = 1$, the optimization will in general stop either when $J[\epsilon_{(K)}] = 1 - \delta$ or when the gradient of the cost functional vanishes: $\nabla J[\epsilon_{(K)}] \simeq 0$.

Optimal control techniques have been applied to a variety of scenarios ranging diverse areas, and has been especially fruitful in quantum chemistry [8, 27] and quantum information related protocols [28–30]. Here, we will focus on a simple but non-trivial model of a controlled quantum system. Let us consider a two-level system described by the Hamiltonian

$$\hat{H}(\epsilon(t)) = \frac{\Delta}{2} \hat{\sigma}_x + \epsilon(t) \hat{\sigma}_z, \quad (2)$$

where σ_x and σ_z are Pauli matrices. External driving $\epsilon(t)$ weights the energy difference between σ_z eigenstates and Δ is usually referred to as the energy gap, since it measures the minimum separation between the eigenenergy branches of $H(\epsilon)$, which occurs at $\epsilon = 0$. The Landau–Zener Hamiltonian (equation (2)) can describe a two-level atom interacting with electromagnetic radiation in the rotating-wave approximation (RWA) for constant Rabi frequency and a time-modulated detuning [31] or a spin in a constant magnetic field in the x -direction and a time-dependent one in the z -direction.

When the state is initially prepared as $|\psi(t \rightarrow -\infty)\rangle = |0\rangle$, choosing $\epsilon(t) = v t$ yields the famous Landau–Zener problem [22, 32], for which an analytical formula can be drawn for the asymptotic probability of population transfer between $|0\rangle$ and $|1\rangle$ (the eigenstates of σ_z). Here, we are interested in achieving complete population transfer between those states, and we will often also be interested in minimizing the evolution time. Linear sweeping of the control parameter is not efficient in this context, since large evolution times (scaling as Δ^{-2} [33]) would be required by virtue of the adiabatic theorem. As a consequence, we will resort to optimization techniques in order to find an appropriate shape for $\epsilon(t)$ that maximizes

$$J[\epsilon] = |\langle 1 | \hat{U}_T[\epsilon] | 0 \rangle|^2 \quad (3)$$

for each fixed value of T . We note that the control landscape could have symmetries depending on the particular protocol that we want to implement. For the case of complete population transfer between σ_z eigenstates, any control protocol has the properties $J[\epsilon(t)] = J[\epsilon(T-t)]$ and $J[\epsilon(t)] = J[-\epsilon(T-t)]$. We give more details about this in the appendix.

The issue of time-optimal control in this scenario was studied by Hegerfeldt [34], who showed that there is a minimum control time which is given by

$$T = T_{\min} = \frac{\pi}{\Delta}, \quad (4)$$

meaning that for $T < T_{\min}$, it is not possible to achieve full population transfer, i.e. $J[\epsilon] < 1$ for all $\epsilon(t)$. Remarkably, the field shape which accomplishes the control task at $T = T_{\min}$ is simply $\epsilon(t) = 0$.

3. Looking at the landscape: toy model for the control

As already mentioned in the introduction, it was shown in [15] that the control landscape $J[\epsilon]$ determined by the Landau–Zener Hamiltonian (2) has only global optima when we consider a continuous-in-time control field $\epsilon(t)$, that is, it is a trap-free model. Moreover, these global optima should correspond to $J[\epsilon] = 1$ only for $T \geq T_{\min}$. However, any practical realization of the optimization problem stated above implies performing a coarse graining of the temporal variable. As a result, the control function $\epsilon(t)$ will now be represented by a vector of control variables, namely

$$\epsilon(t) \rightarrow \{\epsilon_k\} \equiv \vec{\epsilon} \quad (5)$$

with $k = 1, 2, \dots, N_{\text{ts}}$. The functional dependence of J is then mapped to an explicit dependence of the objective on the control parameters (and, also, on the evolution time)

$$J[\epsilon] \rightarrow J(\{\epsilon_k\}, T). \quad (6)$$




Figure 1. Control landscape for the Hamiltonian of equation (2) with initial state $|0\rangle$ and target state $|1\rangle$. The case depicted here corresponds to the scenario where $N_{ts} = 2$ and so the control field is mapped to a 2D vector (a_1, a_2) . Subplots correspond to different values of the total evolution time: (a) $T/T_{\min} = 0.7$, (b) $T/T_{\min} = 1.0$, (c) $T/T_{\min} = 1.2$ and (d) $T/T_{\min} = 10$. The energy gap is set to $\Delta = 1$ in all cases.

Naturally, we will not be able to visualize the landscape for $N_{ts} > 2$, and because of that, in section 4, we will propose various strategies to obtain information about its features. In this section, we will study the simplest non-trivial scenario, where $N_{ts} = 2$, in order to gain intuition about the properties of the landscape. To that end, we propose that the field $\epsilon(t)$ is of the form

$$\epsilon(t) = \begin{cases} a_1 & \text{if } t \leq T/2 \\ a_2 & \text{if } t > T/2. \end{cases} \quad (7)$$

It is then easy to evaluate the objective functional of equation (3):

$$J(a_1, a_2, T) = \left| \langle 1 | e^{-iH(a_2)\frac{T}{2}} e^{-iH(a_1)\frac{T}{2}} | 0 \rangle \right|^2. \quad (8)$$

In figure 1, we plot the control landscape of equation (8), as a function of control parameters (a_1, a_2) for different values of the total evolution time T , both below and above the quantum speed limit time T_{\min} (equation (4)). Figure 1(a) corresponds to $T = 0.7T_{\min}$. In this case, the landscape shows only a single global maximum at the origin, with maximum fidelity of $J \simeq 0.65$ (sub-optimal). As $T \rightarrow T_{\min}$ the landscape's topological structure remains unaltered, hosting a single global maximum with ever-growing fidelity. At $T = T_{\min}$, the global maximum reaches its optimal height $J = 1$ (figure 1(b)).

For $T > T_{\min}$, a much more intricate structure in the topology of $J(a_1, a_2)$ is clearly observed in figures 1(c) and (d). The maximum at the origin splits into two symmetrical global maxima, which steadily separate from each other as the total evolution time is increased figure 1(c). In figure 1(d), we plot the control landscape at $T = 10T_{\min}$. The landscape presents many extrema, as was shown in [15]. We observe a shrinking of the characteristic scales as a function of control parameters (a_1, a_2) for increasing T .

It is important to point out that, in general, the minimum control time will be a function of N_{ts} . This is so because we have defined T_{\min} as the minimum value of T such that $J(\{\epsilon_k\}, T) = 1$. However, for the two-level problem considered here, analytical arguments have shown that $T_{\min} = \pi/\Delta$ [34] and the corresponding optimal field is $\epsilon(t) = 0$. Since that field is trivially achieved with any discretization (i.e. any value of N_{ts}), we do not have to worry here about having to consider a N_{ts} -dependent minimum control time.

As mentioned in section 2, the strategy to obtain control fields through optimization is to propose an initial seed field and update it iteratively by using information about the gradient of $J[\epsilon]$. This process stops either when unit fidelity is obtained ($1 - J < 10^{-16}$) or when the

sum of the squares of the gradients with respect to the control amplitudes falls below a certain threshold ($\sum_i |\nabla_i J|^2 < 10^{-6}$). We can interpret this procedure as a path through the landscape. In [21, 35, 36], a measure R of the straightness of the paths was proposed. R is defined as the ratio between the length of the optimization trajectory, defined as

$$d_{\text{PL}} = \int_0^{s_{\text{max}}} \left[\frac{1}{T} \int_0^T \left(\frac{\partial \epsilon(s, t)}{\partial s} \right)^2 dt \right]^{1/2} ds, \quad (9)$$

and the Euclidean distance between the initial seed and optimal control field,

$$d_{\text{EL}} = \left[\frac{1}{T} \int_0^T [\epsilon(s_{\text{max}}, t) - \epsilon(0, t)]^2 dt \right]^{1/2}. \quad (10)$$

The new variable s parametrizes the optimization path, such that the initial seed is $\epsilon(s = 0, t)$ and the optimized solution is given by $\epsilon(s = s_{\text{max}}, t)$. We stress that the importance of inquiring into the non-topological details of the landscape is evident, since a trap-free landscape does not guarantee an easy optimization. Complex-structured landscapes may constrain optimization paths to inefficient twisted routes.

In figure 2, we show some results of the calculation of R using a simple steepest ascent algorithm. For each initial seed, the step size of the gradient optimizer was successively decreased until a converged value of R was obtained. Two representative trajectories through the landscape of equation (8) for $T = 0.7T_{\text{min}}$ are plotted in figure 2(a). As we can see, one of them is completely straight, giving $R = 1$ whereas the other one, yielding $R = 1.23$, is slightly arched [21]. In order to gain some insight into the structure of the landscape, a statistical analysis of R was performed using 1000 random initial seed fields in the region $-1 < a_1, a_2 < 1$. Each initial field was optimized and the measure R of its path in the landscape was computed. Mean value and standard deviation of the R distributions are shown in figure 2(b) for various values of T . Only those trajectories leading to a global maximum were considered. This figure indicates that the length R of the path towards the optimal, and therefore its complexity, is increased with T/T_{min} . This is consistent with the structure of the landscape that was plotted in figures 1(a)–(d).

4. General control fields and multidimensional landscapes

The control landscape cannot be directly visualized for control space dimensions above $N_{\text{ts}} = 2$. We can, however, obtain information about it by generating a large number of initial seeds and analyzing the resulting control trajectories statistically. This idea has been used in previous works on quantum optimal control [15, 18]. This approach has intrinsic limitations, since limited computational resources give rise to what is usually known as the exploration–exploitation trade-off in optimization theory [37], by means of which a *detailed* characterization the *whole* multidimensional landscape is out of reach. Here, we will focus on probing the landscape in a region centered around $\epsilon(t) = 0$. This is an obvious reference in this case, since the energy spectrum of the Hamiltonian in equation (2) is symmetric with respect to $\epsilon = 0$. Also, using constant, feature-less fields as initial guesses is a common approach in optimal control problems. Finally, as mentioned in section 2, this is the actual optimal field for $T = T_{\text{min}}$. Thus, the chosen landscape region is relevant for exploring.

In the remainder of this section, we propose various methods for probing the structure and topology of the control landscape. The common methodology is as follows. An initial guess for the control field $\epsilon^{(0)}(t)$ is generated as a vector of random numbers $\{\epsilon_k^{(0)}\}$, where




Figure 2. (a) Two optimization trajectories in the landscape of equation (8) for $T/T_{\min} = 0.7$. The corresponding initial seeds are shown as orange circles, and the final optimized parameters are shown as green circles. (b) Mean value of R as a function of T/T_{\min} using 1000 initial seeds. Vertical dashed line indicates the value of T/T_{\min} at which the trajectories of panel (a) were computed. Vertical orange lines represent the standard deviation of the distributions.

$-A \leq \epsilon_k^{(0)} \leq A$ and $k = 1, \dots, N_{\text{ts}}$. The field is then optimized using the GRAPE algorithm, which is currently built into the QuTiP Python package [38, 39]. The iterative optimization stops when the gradient of the functional satisfies a standard convergence criterion. The process is then repeated for a large number of random initial seeds (of the order of 1000) in order to draw sufficient statistics.

4.1. Distance between optimal fields

Our first focus of interest is on the topology of the landscape—specifically, the number, distribution and nature of its extrema. When dealing with multidimensional control optimization, it is a well known fact that different initial guesses lead generally to different optimized fields, albeit usually yielding similar optimized fidelities. In order to explore this behavior, we study the distribution of the optimized control fields as follows. For each pair of optimal fields found, we calculate the distance between them simply as

$$D_{ij} = \frac{1}{T} \int_0^T \left| \epsilon^{(i)}(t') - \epsilon^{(j)}(t') \right| dt'. \quad (11)$$

The mean value of the distance $\langle D \rangle$ between optimized fields is plotted in figure 3 as a function of the evolution time T , and for different values of the number of time slots N_{ts} . There, we can see that $\langle D \rangle$ is close to zero when we intend to control the system below the quantum




Figure 3. Mean value of the distance between optimal fields as defined in equation (11), as a function of the total evolution time measured in units of $T_{\min} = \pi/\Delta$. The results of a statistical analysis, involving one thousand random initial seeds in the $A = 1$ region, are shown for various choices of N_{ts} , the number of time slots into which the field is discretized.

speed limit time, and rises steadily beyond that regime. This tells us that the control landscape for $T/T_{\min} < 1$ has essentially a single maximum to which all initial seeds converge. However, for longer control times, optimized solutions spread out in multiple global maxima. This is the same feature that was found for the simple case of $N_{ts} = 2$, and we show here that it extends to N_{ts} of the order of 1000.

A remarkable conclusion from this analysis is that the landscape undergoes a sudden transformation at $T = T_{\min}$. Note that, *a priori*, we expected the evolution time T to impact the height of the extrema, by definition. Nevertheless, results shown here demonstrate a more profound topological change on the landscape when traversing the quantum speed limit, giving rise to multiple extrema which were absent for smaller control times. It would be interesting to explore whether this phenomenon also takes place in more complex quantum systems. We leave this issue for future work.

We also point out that the spreading of optimal solutions at the onset of controllability, i.e. for $T > T_{\min}$, is consistent with the concept of *superlandscape* introduced recently. This is due to the fact that small (time localized) perturbations of an optimum field can be easily compensated to make the perturbed field optimal as well, leading to many closely spaced locally optimal solutions [37].

In order to obtain a deeper insight about this result, we take a closer look at the actual distribution of distances found by this procedure. Results are shown in figure 4(a) for two different values of the evolution time T . For $T/T_{\min} = 1$, distances between optimized fields form a unimodal distribution, indicating a narrow spread of optimal solutions around some point in parameter space. As already seen in figure 3, for $T/T_{\min} > 1$ the mean distance shifts to larger values. Also interestingly, the distribution becomes bimodal in this case. This means that the optimized solutions now cluster around two points in control space \vec{e}_A and \vec{e}_B ; the leftmost peak in the distribution corresponds to distances between solutions in the same cluster, and the rightmost peak to solutions in different clusters. This clustering behavior can be understood as the emergence of the two global maxima in the superlandscape. This relation between optimal fields was already seen in the 2D case $N_{ts} = 2$ —see figure 1(c). This shows a non-trivial connection between the easily tractable low-dimensional control landscape and the complex




Figure 4. (a) Distribution of distances between control fields as calculated by equation (11), for two different values of the evolution time T . The initial seed fields for optimization were generated as vectors of random numbers in the $A = 1$ region. The number of time slots N_{ts} was fixed to a hundred. Notice that histograms for $T/T_{\min} < 1$ have zero mean and vanishingly small variance, and so are not shown in this plot. Panels (b) and (c) show optimized control fields belonging to different regions (clusters) in control space \vec{e}_A and \vec{e}_B (see text for details).

multi-dimensional one. We plot representative solutions of each type in figures 4(b) and (c), from which we observe that $\vec{e}_A \sim -\vec{e}_B$.

4.2. Trapping probability

In this section we explore the emergence of traps, i.e. sub-optimal local maxima in the landscape, as a function of the constraints imposed on the control problem. In order to do this, we consider that a particular control trajectory has become trapped if it is unable to reach some fixed target fidelity $\tilde{J} < 1$. Then, for fixed T and N_{ts} , we define the trapping probability as the fraction of optimized solutions that became trapped.

In figure 5, we plot this quantity as a function of the number of time slots N_{ts} , for various evolution times T/T_{\min} and for two values of the fidelity threshold $\tilde{J} = 0.99$ (figure 5(a)) and 0.999 (figure 5(b)). It is readily seen from the plots that the trapping probability goes to zero for large N_{ts} in all cases. This had been shown in [15], where the authors analytically proved that in the limit $N_{\text{ts}} \rightarrow \infty$, the landscape is indeed devoid of traps.

From figure 5, it can also be seen that for values of T far from the quantum speed limit, the different curves approach each other, meaning there is no significant change in the abundance of traps with total time variation. However, when T/T_{\min} approaches 1, the behavior is markedly different, and the trapping probability decay slows down. Moreover, this effect becomes more pronounced for a tighter fidelity threshold; that is, for evolution times $T/T_{\min} > 1.3$, the probability of trapping does not change with the choice of threshold. This implies that on the onset of controllability, the optimization becomes harder—as it is more likely for a random initial seed to converge to a trap. This interesting result highlights the role of the evolution time as an important constraint in control problems.

We point out that many previous works have reported that optimal control near the quantum speed limit time tends to become *slower*, in the sense that more iterations are needed in order to reach a satisfactory fidelity. In the landscape picture, this can be understood as maxima

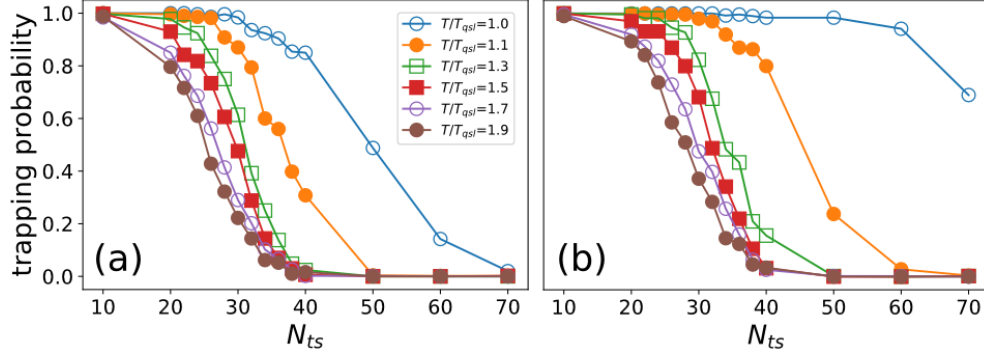


Figure 5. Trapping probability as a function of N_{ts} (the number of parameters in the control field), for different values of the total evolution time T . A control trajectory is said to be trapped if it converges to a final fidelity below (a) 0.99 and (b) 0.999. Results shown correspond to 500 random seeds, in the region of parameter space where $A = 50$ (see text for more details).

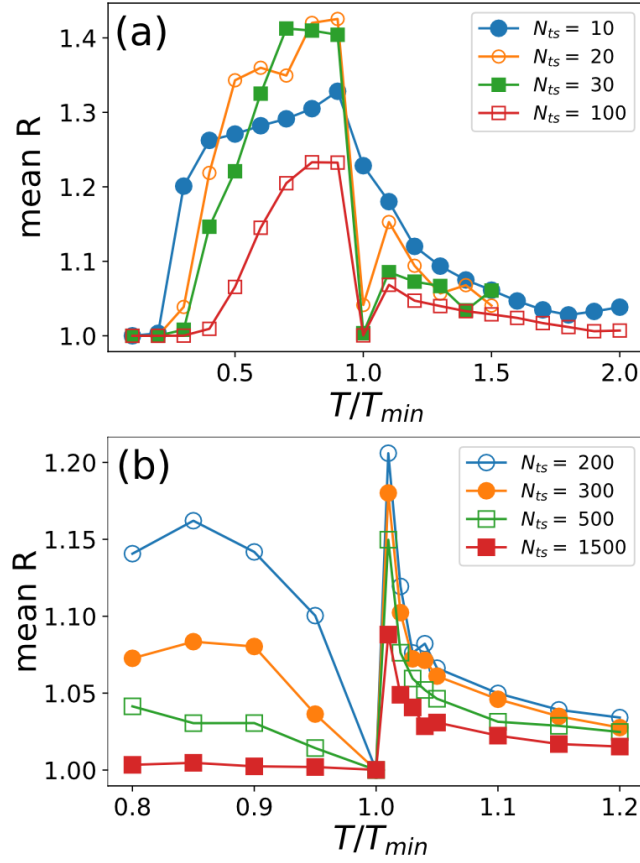


Figure 6. Mean value of R as a function of T/T_{min} for several number of times slots N_{ts} of the control field. We have used 1000 initial seeds. (a) $N_{ts} = 10, 20, 30$ and 100. (b) $N_{ts} = 200, 300, 500$ and 1000.

becoming more *flat* (as can be seen for example in figure 1(b)). We stress that the result shown in this work is of a different nature, since here we observe the appearance of traps for sufficiently small times. Since we set a convergence criterion for the optimization based on the gradient of the functional, we can effectively distinguish actual traps from flat global optima.

4.3. Looking at the structure of the landscape using the measure R

In this section, we study the behavior of the measure R introduced previously, in the case where $N_{ts} > 2$ so neither the landscape nor optimization paths can be directly visualized. Instead, we probe the structural (non-topological) features of the landscape by analyzing the behavior of the R distributions.

In figure 6, we show the mean value of R computed for 1000 initial random seeds for several values of time slots N_{ts} of the field. Several important features can be remarked. For small $T/T_{min} \rightarrow 0$ in figure 6(a), we can see that the mean R approaches 1. This means that almost all trajectories of the optimization process are straight lines, showing that the landscape has a simple topological structure as was shown for $N_{ts} = 2$ in figure 1(a). Remarkably, the same is observed for $T = T_{min}$ and for $N_{ts} > 10$ (see the pronounced deep trench in $T/T_{min} = 1$). That is, the landscape has also a simple structure in T_{min} . For greater optimal evolution time T the mean R shows a sharp jump that indicates the birth of the two optimal solutions shown in figure 4. We note that as T is increased after the T_{min} , the mean value of R decays smoothly. This fact shows that the structure of the landscape is simpler when the constraint on the evolution time T is relaxed, as expected [9, 21].

In figure 6(b), we show the mean value of R for greater values of time slots $N_{ts} = 200, 300, 500$ and 1000. We can clearly see that the greater is N_{ts} , the smaller is R . This means that the optimization process becomes simpler as we increase the number of time slots into which the field is divided.

5. Concluding remarks

In this paper, we study the quantum control landscape of a paradigmatic system: the two-dimensional Landau–Zener Hamiltonian. We present systematic study of the influence of two important constraints: the discretization and the time extension of the control field. When the number of time steps of the control field is two, the landscape can be directly plotted and several interesting features are easily visualized. In particular, we find that the QCL in this case undergoes a profound transformation at the quantum speed limit, in which extrema split up to make way for multiple optimal solutions.

For control fields with more than two time slots, we use indirect methods to unravel the structure of the landscape. We consider the distances between optimal fields, which allows us to map the topological structure of the landscape (for e.g. number and distribution of maxima). Regarding the fidelities of those optimal fields, we study the emergence of traps due to the constraints imposed to the problem, and show that for a fixed discretization, landscapes with T closer to the T_{min} are more likely to show traps. We also compute the R metric defined in [21], which is a measure of the straightness of the path taken by the optimizer through the landscape, connecting an initial seed with the corresponding optimized solution. This measure gives us information about the structural properties of the landscape. These methods allow us to characterize the QCL in these cases, and we believe that results shown here serve as a testbed for future studies in more complex quantum systems. An interesting byproduct of our study is that many of the features directly visualized in the 2D landscape carry over to higher

dimensions (with N_{ts} as large as 1000). It would be interesting to analyze to what extent we can learn about the QCL by using a small-dimensional ansatz for the control field in other types of quantum system, a matter which we leave for future study.

Considering that the quantum control landscape contains the relevant information for coherent control, and that we have been able to unravel its structure for a simple system, our work represents a step forward to the understanding of this important functional in more complex systems, especially its behavior near the minimum control time. We also note that in [40], the optimization complexity for the same model has been studied using Reinforcement Learning techniques, and many connections with our results can be observed. Altogether, this deeper understanding of the control properties paves the way for the systematic generation of high-speed protocols that can effectively control real quantum systems facing decoherence.

Acknowledgments

PMP acknowledges interesting discussions with Jens-Jakob Sorensen during the GRC Conference of Quantum Control in South Hadley. We would also like to thank Alexander Pitchfork of the QuTiP team and Gabriele De Chiara for the fruitful insights. This work was partially supported by CONICET Grant No. PIP 112201 501004 93CO, UBACyT Grant No. 20020130100406BA and National Science Foundation Grant No. PHY-1630114.

Appendix. Symmetries in the QCL

Consider a control task where the final state $|f\rangle$ is obtained by rotating the initial state $|i\rangle$ an angle π about y-axis,

$$e^{\frac{-i\pi\hat{\sigma}_y}{2}}|i\rangle = |f\rangle. \quad (\text{A.1})$$

The Landau–Zener propagator (equation (1)) is invariant under such rotation; thus, any control task given by (A.1) has an associated QCL that exhibits the following symmetry:

$$\begin{aligned} J[\epsilon(t)] &= |\langle f|\hat{U}_T[\epsilon(t)]|i\rangle|^2 \\ &= |\langle i|e^{\frac{-i\pi\hat{\sigma}_y}{2}}\hat{U}_T[\epsilon(t)]e^{\frac{i\pi\hat{\sigma}_y}{2}}|f\rangle|^2 \\ &= |\langle i|U_T[\epsilon(t)]|f\rangle|^2 \\ &= J[\epsilon(T-t)]. \end{aligned}$$

If we consider a control task with initial and final states related by a π rotation about x-axis, it is easy to show that $J[\epsilon(t)] = -J[\epsilon(T-t)]$.

ORCID iDs

Pablo M Poggi  <https://orcid.org/0000-0002-9035-3090>

References

- [1] Nielsen M A and Chuang I L 2000 *Quantum Computation and Quantum Information* (Cambridge: Cambridge University Press)
- [2] Gisin N and Thew R 2007 Quantum communication *Nat. photon.* **1** 165

- [3] Georgescu I M, Ashhab S and Nori F 2014 Quantum simulation *Rev. Mod. Phys.* **86** 153
- [4] Barends R *et al* 2016 Digitized adiabatic quantum computing with a superconducting circuit *Nature* **534** 222
- [5] Yin J *et al* 2017 Satellite-based entanglement distribution over 1200 kilometers *Science* **356** 1140–4
- [6] Bernien H *et al* 2017 Probing many-body dynamics on a 51-atom quantum simulator *Nature* **551** 579
- [7] Peirce A P, Dahleh M A and Rabitz H 1988 Optimal control of quantum-mechanical systems: existence, numerical approximation, and applications *Phys. Rev. A* **37** 4950
- [8] Somló J, Kazakov V A and Tannor D J 1993 Controlled dissociation of I_2 via optical transitions between the x and b electronic states *Chem. Phys.* **172** 85–98
- [9] Rabitz H A, Hsieh M M and Rosenthal C M 2004 Quantum optimally controlled transition landscapes *Science* **303** 1998–2001
- [10] Brockett R 1991 Dynamical systems that sort lists, diagonalize matrices, and solve linear programming problems *Linear Algebra Appl.* **146** 79–91
- [11] Shen Z, Hsieh M and Rabitz H 2006 Quantum optimal control: Hessian analysis of the control landscape *J. Chem. Phys.* **124** 204106
- [12] Hsieh M, Wu R and Rabitz H 2009 Topology of the quantum control landscape for observables *J. Chem. Phys.* **130** 104109
- [13] Moore K W and Rabitz H 2012 Exploring constrained quantum control landscapes *J. Chem. Phys.* **137** 134113
- [14] Pechen A N and Tannor D J 2012 Quantum control landscape for a λ -atom in the vicinity of second-order traps *Isr. J. Chem.* **52** 467–72
- [15] Pechen A and Il'in N 2012 Trap-free manipulation in the Landau–Zener system *Phys. Rev. A* **86** 052117
- [16] Zhdanov D V and Seideman T 2015 Role of control constraints in quantum optimal control *Phys. Rev. A* **92** 052109
- [17] Zhdanov D V 2017 Theory of quantum control landscapes: overlooked hidden cracks (arXiv:1710.07753)
- [18] Sørensen J J *et al* 2016 Exploring the quantum speed limit with computer games *Nature* **532** 210
- [19] Gajdacz M, Das K K, Arlt J, Sherson J F and Opatrny T 2015 Time-limited optimal dynamics beyond the quantum speed limit *Phys. Rev. A* **92** 062106
- [20] Tibbetts K W M, Brif C, Grace M D, Donovan A, Hocker D L, Ho T-S, Wu R-B and Rabitz H 2012 Exploring the tradeoff between fidelity and time optimal control of quantum unitary transformations *Phys. Rev. A* **86** 062309
- [21] Nandori A, Donovan A, Ho T-S and Rabitz H 2013 Exploring quantum control landscape structure *Phys. Rev. A* **88** 033425
- [22] Zener C 1932 Non-adiabatic crossing of energy levels *Proc. R. Soc. A* **137** 696
- [23] Shevchenko S N, Ashhab S and Nori F 2010 Landau–Zener–Stückelberg interferometry *Phys. Rep.* **492** 1–30
- [24] Zurek W H, Dorner U and Zoller P 2005 Dynamics of a quantum phase transition *Phys. Rev. Lett.* **95** 105701
- [25] Konnov A I and Krotov V F 1999 On global methods for the successive improvement of control processes *Avtom. Telemekh.* **60** 77–88
- [26] Khaneja N, Reiss T, Kehlet C, Schulte-Herbrüggen T and Glaser S J 2005 Optimal control of coupled spin dynamics: design of nmr pulse sequences by gradient ascent algorithms *J. Magn. Reson.* **172** 296–305
- [27] Shi S, Woody A and Rabitz H 1988 Optimal control of selective vibrational excitation in harmonic linear chain molecules *J. Chem. Phys.* **88** 6870–83
- [28] Poulsen U V, Sklarz S, Tannor D and Calarco T 2010 Correcting errors in a quantum gate with pushed ions via optimal control *Phys. Rev. A* **82** 012339
- [29] Goerz M H, Calarco T and Koch C P 2011 The quantum speed limit of optimal controlled phasegates for trapped neutral atoms *J. Phys. B: At. Mol. Opt. Phys.* **44** 154011
- [30] Cross A W and Gambetta J M 2015 Optimized pulse shapes for a resonator-induced phase gate *Phys. Rev. A* **91** 032325
- [31] Shore B W 2011 *Manipulating Quantum Structures Using Laser Pulses* (Cambridge: Cambridge University Press)
- [32] Landau L D 1932 *Phys. Z. Sowjetunion* **2** 46

- [33] Poggi P M, Lombardo F C and Wisniacki D A 2013 Controlling open quantum systems using fast transitions *Phys. Rev. A* **87** 022315
- [34] Hegerfeldt G C 2013 Driving at the quantum speed limit: optimal control of a two-level system *Phys. Rev. Lett.* **111** 260501
- [35] Nanduri A, Ho T-S and Rabitz H 2016 Quantum-control-landscape structure viewed along straight paths through the space of control fields *Phys. Rev. A* **93** 023427
- [36] Nanduri A, Shir O M, Donovan A, Ho T-S and Rabitz H 2015 Exploring the complexity of quantum control optimization trajectories *Phys. Chem. Chem. Phys.* **17** 334–47
- [37] Heck R *et al* 2017 Do physicists stop searches too early? A remote-science, optimization landscape investigation (arXiv:[1709.02230](#))
- [38] Johansson J R, Nation P D and Nori F 2013 Qutip 2: a python framework for the dynamics of open quantum systems *Comput. Phys. Commun.* **184** 1234–40
- [39] Johansson J R, Nation P D and Nori F 2012 Qutip: an open-source python framework for the dynamics of open quantum systems *Comput. Phys. Commun.* **183** 1760–72
- [40] Bukov M, Day A G R, Sels D, Weinberg P, Polkovnikov A and Mehta P 2017 Machine learning meets quantum state preparation. The phase diagram of quantum control (arXiv:[1705.00565](#))